

AUTOMATED SELECTION OF APPROPRIATE INFORMATION  
BASED ON A COMPUTER USER'S CONTEXT

CROSS REFERENCE TO RELATED APPLICATIONS

This application is a continuation-in-part of U.S. Patent Application No. 09/216,193 (Attorney Docket No. 294438001US), filed December 18, 1998 and entitled "Method and System for Controlling Presentation of Information to a User Based on the User's Condition"; of U.S. Patent Application No. 09/464,659 (Attorney Docket No. 294438003US), filed December 15, 1999 and entitled "Storing and Recalling Information to Augment Human Memories"; of U.S. Patent Application No. \_\_\_\_\_ (Attorney Docket No. 294438008US1), filed April 2, 2001 and entitled "Thematic Response To A Computer User's Context, Such As By A Wearable Personal Computer", which claims the benefit of provisional U.S. Patent Application No. 60/193,999 filed April 2, 2000; and of U.S. Patent Application Serial No. \_\_\_\_\_ (Attorney Docket No. TG1-004US), filed April 2, 2001 and entitled "Soliciting Information Based On A Computer User's Context", which claims the benefit of provisional U.S. Patent Applications Nos. 60/194,000 and 60/194,758, both filed April 2, 2000. Each of these applications is hereby incorporated by reference in their entirety.

This application also claims the benefit of provisional U.S. Patent Application No. 60/196,625 (Attorney Docket No. 294438025US), filed April 12, 2000 and entitled "Enhanced Gaming With Wearable Computers", and of provisional U.S. Patent Application No. 60/240,694 (Attorney Docket No. TG1-013), filed October 16, 2000 and entitled "Method For Synchronizing Computer Generated Images With Real World Images". Each of these applications is hereby incorporated by reference in their entirety.

## TECHNICAL FIELD

The present invention relates generally to computer-supported filtering of information, and more particularly to presentation of filtered information (such as unsolicited advertisements) to a computer user based on the user's current context.

## 5 BACKGROUND

As computers become increasingly powerful and ubiquitous, users increasingly employ their computers for a broad variety of tasks. For example, in addition to traditional activities such as running word processing and database applications, users increasingly rely on their computers as an integral part of their daily  
10 lives. Programs to schedule activities, generate reminders, and provide rapid communication capabilities are becoming increasingly popular. Moreover, computers are increasingly present during virtually all of a person's daily activities. For example, hand-held computer organizers (*e.g.*, PDAs) are more common, and communication devices such as portable phones are increasingly incorporating computer capabilities. Thus, users may be presented with output information from one or more computers at any time.

Accompanying the increasing use and portability of computers is an increasing desire to provide information to users through wireless and other communication channels. Advertisers are particularly interested in reaching computer users as such users often represent desirable target markets. Advertisers would like to  
20 provide advertisements to various computer and communication devices that are relevant to the users. Other groups in addition to advertisers may also want to disseminate information to computer users.

With the convergence of computerization, wireless capabilities, digitalization, and ease of dissemination, the amount of potential information that may be  
25 bombarded on individual users may prove too overwhelming. The sheer volume of information may prevent or discourage computer users from making any effort to examine the information and find what is immediately desirable or necessary. However,

simply reducing the amount of information, or the number of information sources, is not an appropriate solution because the user loses access to information of potential interest.

In general, current solutions for selecting incoming information fail because they do not address the dynamic, immediate, and arbitrary desires and needs of a user.

5 The solutions cannot be easily modified, extended, or shared by the user.

One possible solution is to simply ignore some or all available information, such as unsolicited information that is received. This solution is not always possible, however, and is often difficult and inconvenient. For example, it is difficult to ignore an audio message of moderate volume, especially when those messages are specifically  
10 designed to be difficult to ignore (*e.g.*, by compressing the dynamic range so that the average volume is higher). Plus, rejecting all information robs the user of potentially important information in which the user may be interested.

Another solution is to select a set of information sources from which information (solicited and/or unsolicited) is permitted. This exclusionary practice is not  
15 fully desirable because the user loses access to potentially desirable information from a source that may not be on the acceptable source list. Also, even with a small number of information sources, the user's ability to give continuous attention to multiple, or even a single information source, is limited. This is especially true if the user is to remain "in task" and not be distracted by inappropriate information. Junk mail filters used by  
20 conventional email programs attempt to filter unwanted bulk mailings by using keywords or user enumeration. These filters have the drawback in that the user must manually update desired or undesired sources to account for new sources.

Another drawback of conventional solutions is that the solutions are based on the underlying assumption that the user is stationary and in a static context, such as  
25 sitting at a desktop computer or sitting in front of a web-enabled television. The solutions do not take into account that the user's context may be dynamic and changing or evolving over time. Thus, the present-day solutions for filtering digital information are not easily ported to the portable computing environment such as portable computing devices, PDAs, communications devices, and wearable computing devices.

Users, on the other hand, may not mind receiving even unsolicited information (*e.g.*, advertisements) if the information is received at a convenient time and made available in a non-intrusive manner. Much of this timing and convenience is dictated by the user's current context. The user's context can either limit or enhance their ability and desire to receive and process information. When the user is not interested in receiving a message, or the user is unable to receive a specific message because they are too busy, or the user is unable to accept a specific message due to the user's present environment or situation (such as the user's current information presentation system configuration), the context may simply dictate that the user is not interested or able to review any information.

Accordingly, there is a need for improved techniques that enable received information, including unsolicited information, to be received and presented in a timely and relevant manner to the computer user. Such techniques would be desirable for the user because the user will be made aware of potentially relevant information and beneficial for the information provider (such as an advertiser) because the user may not be as apt to simply ignore or discard the information. Instead, the user may want to experience (*e.g.* see, hear, smell, etc.) the advertisement when the user's context is suitable for reviewing the advertisement.

In addition to the general need for improved techniques to select appropriate information and present it in an appropriate manner, there are various specialized environments or situations which harbor additional problems or constraints. For example, existing computer games don't provide a complete gaming environment that reflects the user's detailed context, such as can be maintained or modeled by the user's characterization module. Thus, there is a need for computer games to provide enhanced functionality by using such modeled context information, such as to select appropriate game-related information or functionality based on the user's context and to present it in an appropriate manner.

In addition, in the area of computer-based augmented reality, there is a need for improved techniques to select appropriate information and present it in an appropriate manner. In particular, current technology is unable to dynamically augment reality by

integrating virtual information (such as images, sounds, and other related information) with real world information (*e.g.*, visible objects and audible sounds) without hardware clues. Maintaining an overlay, or visual mapping, of virtual objects on top of real-world objects is known as registration. For example, in an industrial setting an example of registration may involve superimposing a virtual image of a wire over the real-world view of the wire, and having that virtual wire remain properly associated with the real wire while the user's head is moving. Such registration can be facilitated by adding visual clues to the relevant environment, such as patterns of dots on a circuit board containing the wire. However, the need for such visual cues severely limits the usability of augmented reality. Another difficulty in managing the overlay of virtual information onto the real world involves presenting and formatting computer interactions controls without detracting from the user's experiencing of the real world.

## BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 illustrates a system in which messages are pushed from multiple information sources to multiple clients over a network.

Figure 2 illustrates one exemplary implementation of a client, which is embodied as a wearable computer.

Figure 3 illustrates an architecture implemented at the client to filter incoming messages according to the user's context.

Figure 4 illustrates a composite filter consisting of a hierarchy of filters for filtering incoming messages.

Figure 5 is a flow diagram of a process for filtering incoming messages and deciding when to present the messages based on the user's context.

Figure 6 illustrates a message distribution system in which an intermediary entity matches messages to intended users and distributes the messages on behalf of information sources.

Figure 7 is a block diagram illustrating an architecture for augmenting reality by presenting appropriate information based on a user's context.

Figure 8 is a block diagram illustrating additional details of the architecture for augmenting reality.

Figure 9 is a block diagram illustrating an architecture for enhancing gaming environments by providing appropriate information and functionality based on a user's context.

Figure 10 is a flow diagram of an embodiment of a Gaming routine.

Figure 11 is a flow diagram of an embodiment of an Augmented Reality routine.

## DETAILED DESCRIPTION

This disclosure pertains to techniques for filtering received information based on the user's current context in order to determine whether the information is appropriate for presentation to the user. In some embodiments, the received information is unsolicited information, such as advertisements. Using an explicit, dynamic, extensible software model of a user's context, such as one that represents their desires, needs, physical condition, surrounding environment, computing capabilities, and data receiving capabilities, information that is appropriate for the user can be determined. In some embodiments, additional determinations are made as to how, when and/or where such appropriate information should be presented, and the information can then be presented in the appropriate determined manner.

In order to use context information about the user to make appropriate determinations, such context information must be received or generated. In some embodiments, various environmental aspects (*e.g.*, mental, physical, computing, data, etc.) are monitored, and a current context of the user is modeled. In other embodiments, such information is received from other sources that perform the monitoring and/or the modeling. Such context information can be represented in various forms, such as with a variety of context attributes (or "condition variables") that each model a single aspect of the context (*e.g.*, ambient temperature, current latitude, or current user task). In some embodiments, context attributes are used that each have a name and at least one value, and can additionally have other properties such as uncertainty, units and timestamp.

When the context information is available, it can then be used to define a set of one or more filters to be selectively applied to incoming information, such as messages. Each filter can have different filter criteria, and a composite of multiple filters can also be used. Since the user's context is typically non-static and continually evolving, the set of filters currently in use can be updated as the context changes.

Information that satisfies the filter criteria of the current filters is determined to be currently appropriate for presentation, although current presentation of such information may not be appropriate (*e.g.*, current sports news may be appropriate information for a particular user, but presentation of such information may not be appropriate while the user is driving or in a business meeting). Thus, in some embodiments, information that is determined to be appropriate based on the current filters is then evaluated to determine whether it should be presented to the user immediately or instead stored for later presentation. In addition, determinations can be made such as where (*e.g.*, on what device) and how (*e.g.*, via which user senses or with what degree of prominence) the information should be presented. As with selection of filters, the user's context can be used in making the determinations of when, where and how to present the content.

### **Exemplary System**

Fig. 1 shows a system 100 in which multiple information sources 102(1), 102(2), ... 102(N) transmit information over one or more networks 104 to multiple clients 106(1), 106(2), ... 106(M). Information that is unsolicited by the clients is said to be "pushed" from the information sources 102 to the clients 106, while information that is solicited or requested is said to be "pulled" by the clients. While it is particularly important to analyze unsolicited information for appropriateness before presentation (since the client or user to whom the information is being pushed has typically not previously indicated any level of interest in the information), it can also be beneficial to analyze solicited information for appropriateness since the actual information that is received may be different than what the user thought they were requesting and/or since it may still be useful to determine an appropriate manner (*e.g.*, when, where and how) in

which to present the information. While most of the techniques described below are generally applicable to both solicited and unsolicited information, some disclosed techniques emphasize the analyzing of unsolicited information. Conversely, additional techniques that focus primarily on active requesting of information from sources and  
5 corresponding analyzing of such information can be found in co-pending U.S. Patent Application Serial Number \_\_\_\_\_ (Attorney Docket No. TG1-004US), filed April 2, 2001 and entitled "Soliciting Information Based On A Computer User's Context".

The information sources 102 may be implemented in a number of ways, such as a host server at a Website, an email service, a paging service, a financial quote  
10 service, a cellular communications host, a radio transmission unit, and the like. As one exemplary implementation, an information source, as represented by source 102(1), may include a content store 110 to store the information and a content server 112 to serve the content to the clients 106. The information communicated from the information sources may be in any data type (*e.g.*, text, graphics, audio, software executables, video, etc.) and  
15 contain essentially any type of subject matter. Examples of possible information that may be pushed to the clients include streaming media, safety bulletins, future or current event announcements, environment status (weather, traffic), local vicinity status (elevator wait time), object status (where is my bus?), executable code, object descriptors (any computer-compatible enumeration or description of sensors or peripheral devices). As  
20 one particular example, the information may be in the form of unsolicited advertisements pushed to the clients 106 from advertisers.

The network 104 is representative of many different network types, including public networks (*e.g.*, the Internet) and proprietary networks. The network may be implemented using wireless technologies (*e.g.*, RF, microwave, cellular, etc.), wire-  
25 based technologies (*e.g.*, cable, fiber optic, wire, etc.), or any combination. Any one of many diverse protocols and formats may be used to package data and transmit it from the source 102 to the client 106.

The clients 106 may be implemented in a variety of ways, including as computers, portable digital assistants (PDAs), communication devices, and the like. The  
30 clients are equipped with conventional mechanisms to receive the information from the



network 104, such as ports, network cards, receivers, modems, and so on. As described in greater detail below, in some embodiments some or all of the clients may be wearable computers.

Each client in the illustrated embodiment, as represented by client 106(1), is equipped with a Condition-Dependent Output Supplier (CDOS) system 120 that monitors the user and the user's environment and that creates and maintains a model (not shown) of the user's context (or "condition"). As the user moves about in different environments, the CDOS system receives various input information, including explicit user input, sensed user information, and sensed environment information. In the illustrated embodiment, one or more sensors 122 provide data to the CDOS system pertaining to the user's environment.

Each client 106 is also equipped with a context-based information filtering system 124 that utilizes one or more filters 126 to selectively pass/block information from the information sources 102 based on the user's present context. Depending on many factors of the user's current context, the context-based information filtering system 124 may allow certain received information, such as important messages from a spouse or family member, to pass through while concurrently blocking other information. Depending on the user's context, the information determined to be appropriate may then be presented immediately or instead stored for delayed presentation.

### **Exemplary Client**

Fig. 2 illustrates one suitable implementation of client 106 as a body-mounted wearable computer worn by a user 202. The computer 106 includes a variety of body-worn input devices, such as a microphone 210, a hand-held flat panel display 212 with character recognition capabilities, and various other user input devices 214. Examples of other types of input devices with which a user can supply information to the computer 106 include speech recognition devices, traditional qwerty keyboards, chording keyboards, half qwerty keyboards, dual forearm keyboards, chest mounted keyboards, handwriting recognition and digital ink devices, a mouse, a track pad, a digital stylus, a

finger or glove device to capture user movement, pupil tracking devices, a gyropoint, a trackball, a voice grid device, digital cameras (still and motion), and so forth.

The computer 106 also has a variety of body-worn output devices, including the hand-held flat panel display 212, an earpiece speaker 216, and a head-mounted display in the form of an eyeglass-mounted display 218. Other output devices 220 may also be incorporated into the computer 106, such as a tactile display, other tactile output devices, an olfactory output device, and the like.

The computer 106 may also be equipped with one or more various body-worn user sensor devices 222. For example, a variety of sensors can provide information about the current physiological state of the user and about current user activities. Examples of such sensors include thermometers, sphygmometers, heart rate sensors, shiver response sensors, skin galvanometry sensors, eyelid blink sensors, pupil dilation detection sensors, EEG and EKG sensors, sensors to detect brow furrowing, blood sugar monitors, etc. In addition, sensors elsewhere in the near environment can provide information about the user, such as motion detector sensors (*e.g.*, whether the user is present and is moving), badge readers, still and video cameras (including low light, infrared, and x-ray), remote microphones, etc. These sensors can be both passive (*i.e.*, detecting information generated external to the sensor, such as a heart beat) and active (*i.e.*, generating a signal to obtain information, such as sonar or x-rays).

The computer 106 may also be equipped with various environment sensor devices 224 that sense conditions of the environment surrounding the user. For example, devices such as microphones or motion sensors may be able to detect whether there are other people near the user and whether the user is interacting with those people. Sensors can also detect environmental conditions that may affect the user, such as air thermometers or Geiger counters. Sensors, either body-mounted or remote, can also provide information related to a wide variety of user and environment factors including location, orientation, speed, direction, distance, and proximity to other locations (*e.g.*, GPS and differential GPS devices, orientation tracking devices, gyroscopes, altimeters, accelerometers, anemometers, pedometers, compasses, laser or optical range finders, depth gauges, sonar, etc.). Identity and informational sensors (*e.g.*, bar code readers,

biometric scanners, laser scanners, OCR, badge readers, etc.) and remote sensors (*e.g.*, home or car alarm systems, remote camera, national weather service web page, a baby monitor, traffic sensors, etc.) can also provide relevant environment information.

The computer 106 further includes a central computing unit 230. The various inputs, outputs, and sensors are connected to the central computing unit 230 via one or more data communications interfaces 232 that may be implemented using wire-based technologies (*e.g.*, wires, coax, fiber optic, etc.) or wireless technologies (*e.g.*, RF, etc.). A receiver 234 is designed to receive incoming messages from the network 104 and to pass them to the central computing unit 230 via the data communications interfaces 232. The receiver 234 may be implemented in a number of ways, including as an RF receiver, a microwave receiver, a satellite receiver, a network card, a wireless network card, a modem, and so forth.

The central computing unit 230 includes a central processing unit (CPU) 240, a memory 242, and a storage device 244. The memory 242 may be implemented using both volatile and non-volatile memory, such as RAM, ROM, Flash, EEPROM, disk, and so forth. The storage device 244 is typically implemented using non-volatile permanent memory, such as ROM, EEPROM, diskette, memory cards, and the like.

One or more application programs 246 can be optionally stored in memory 242 and executed by the CPU 240. These application programs generate information that may be presented to the user via one or more of the output devices.

In the illustrated implementation, a CDOS system 120 is shown stored in memory 242 and executing on the processing unit 240. The CDOS system 120 monitors the user and the user's environment, and creates and maintains an updated model of the current context of the user. As the user moves through different environments, the CDOS system receives various input information including explicit user input, sensed user information, and sensed environment information. The CDOS system updates the current model of the user condition, and presents output information to the user via appropriate output devices. A more detailed explanation of the CDOS system 120 may be found in co-pending U.S. Patent Application Serial No. 09/216,193, and a more detailed

discussion of context attributes (or condition variables) may be found in co-pending U.S. Patent Application Serial No. 09/724,902.

The context-based information filtering system 124 is also shown stored in memory 242 and executing on the processing unit 240. It utilizes data from the modeled user context (*e.g.*, via the CDOS system 120) to selectively filter information according to the user's current environment in order to determine whether the information is appropriate for presentation. The filtering system 124 employs one or more filters 126 to filter the information. The filters may be pre-constructed and stored for subsequent utilization when conditions warrant their particular use, or alternatively the filtering system 124 may construct the filters dynamically as the user's context evolves. In addition, in some embodiments each filter is stored as a distinct data structure (optionally with associated logic), while in other embodiments filters can be provided as logic based on the current context, such as one or more interacting rules provided by the filtering system or CDOS system.

The body-mounted computer 106 may be connected to one or more networks through wired or wireless communication technologies (*e.g.*, wireless RF, a cellular phone or modem, infrared, physical cable, a docking station, etc.). For example, the body-mounted computer of a user could make use of output devices in a smart room, such as a television and stereo when the user is at home, if the body-mounted computer can transmit information to those devices via a wireless medium or if a cabled or docking mechanism is available to transmit the information. Alternately, kiosks or other information devices can be installed at various locations (*e.g.*, in airports or at tourist spots) to transmit relevant (and typically, unsolicited) information to body-mounted computers within the range of the information device.

It is noted that the body-mounted computer 106 is merely one example of a suitable client computer. There are many other implementations of client computing devices that may be used to implement the context-based information filtering system. In addition, while the computer 106 is illustrated in Figure 2 as containing various computing and storage resources and associated I/O devices, in other embodiments a computer 106 may act as a thin client device that receives some or all of its computing

and/or storage capabilities from a remote server. Such a thin client device could consist only of one or more I/O devices coupled with a communications mechanism with which to interact with the remote server.

Thus, those skilled in the art will appreciate that computer systems 106 are merely illustrative and are not intended to limit the scope of the present invention. Such computer systems may be connected to other devices that are not illustrated, including through one or more networks such as the Internet or via the World Wide Web (WWW). In addition, the functionality provided by the illustrated components may in some embodiments be combined in fewer components or distributed in additional components. Similarly, in some embodiments the functionality of some of the illustrated components may not be provided and/or other additional functionality may be available.

Those skilled in the art will also appreciate that, while items (*e.g.*, filters and software components) are illustrated as being stored in memory while being used, these items or portions of them can be transferred between memory and other storage devices for purposes of memory management and data integrity. Similarly, items illustrated as being present on storage while being used can instead be present in memory and transferred between storage and memory. Alternately, in other embodiments some or all of the software modules may execute in memory on another device. Some or all of the described components or data structures may also be stored (*e.g.*, as instructions or structured data) on a computer-readable medium (*e.g.*, a hard disk, a memory, a network, or a portable article to be read by an appropriate drive), and can be transmitted as generated data signals (*e.g.*, as part of a carrier wave) on a variety of computer-readable transmission mediums (*e.g.*, wireless-based and wired/cable-based mediums). In addition, a “client” or “server” computing device may comprise any combination of hardware or software that can interact, including computers, network devices, internet appliances, PDAs, wireless phones, pagers, electronic organizers, television-based systems and various other consumer products that include inter-communication capabilities. Accordingly, the present invention may be practiced with other computer system configurations.

### Context-Based Information Filtering System

Fig. 3 illustrates one exemplary architecture 300 in which the context-based information filtering system 124 may be implemented. The architecture includes several functional components that filter information (*e.g.*, unsolicited information such as advertisements) based on the user's context before allowing that information to be presented to the user.

According to architecture 300, incoming messages are received at a Receiver 234 from a communications network. In some embodiments, the Receiver is a hardware device, while in other embodiments the Receiver could be an application program such as a Web browser. There are a wide variety of messages originating from a number of different sources, including directed messages like telephone and fax calls, email messages, instant messaging, broadcast messages like TV and radio, multicast or streaming messages, etc. Messages can also be computer generated (*e.g.*, from an application program), either remotely or locally. While the illustrated embodiment discusses the analysis of incoming messages, other information that is less structured than messages can also be filtered in a similar manner, such as ambient conversations or visual information of an objectionable nature.

The context-based information filtering system 124 applies one or more filters 126 to incoming messages 302. The filters 126 may be designed and stored as separate from the filtering system or instead integrated within the system. Moreover, the filters may reside at the client or instead be remote from the client, such as in the distributed system described below with respect to Fig. 6. Individual filters may be static in that they are pre-configured for particular usage and stored in the client's computer memory for retrieval by the filtering system 124. Alternatively, individual filters may be generated dynamically by the filtering system 124 as user conditions dictate, such as by applying logic associated with the filtering system or the CDOS system. For instance, if the user is currently sleeping and the computer system detects the user is in a sleep state, the filtering system 124 may generate a temporary filter that blocks all but the highest important incoming messages. Those skilled in the art will appreciate that such a

temporary filter may not be stored as a distinct data structure, but instead reflect processing logic/functionality provided by one of the components.

The filters are designed with two competing goals in mind: (1) precision—receiving only messages that are appropriate to the desire, situation and needs of the user, and (2) completeness—making sure that nothing of interest is missed. The filters 126 can also be general or very specific. For example, one filter may be configured to accept advertisements from soft-drink advertisers, while another filter may be configured to block diet soft-drink advertisements. More specifically, a filter may accept advertisements for a specific cola, but prevent advertisements from a competitive cola product. Thus, if the user only drinks Coca-Cola® brand cola, an advertiser may pay for access to users who never purchase a Pepsi-Cola® brand cola and those constraints could be configured into a specific filter. That filter may also specify an incentive level that an advertiser of a Pepsi-Cola® brand product may need to pay or otherwise provide in order to purchase the privilege of presenting a Pepsi-Cola® brand advertisement to the user who typically only drinks Coca-Cola® brand cola. The incentive level may even be variable for different users. As a way to purchase that privilege, the advertising company may provide coupons to the user or pay the user for simply viewing the advertisement.

Messages that survive the filters 126 are either presented to the user at that time or stored in the client computer for delayed presentation. This is represented by accepted message 304 that has survived the filters 124 and is being forwarded for the user's attention.

The illustrated embodiment of the context-based information filtering system 124 includes a characterization module 310 that attempts to characterize the user's current context and apply the appropriate filters 126 to the incoming messages 302 received at receiver 234. The characterization module 310 also provides information about the user's current context that facilitates the determination of whether the acceptable message(s) should be stored or presented immediately to the user. Whether the messages received at receiver 234 are first passed to the characterization module 310 or to the filters 126 is a matter of design choice. The messages may be first passed to the characterization module 310 and then subjected to the various filters 126 (as shown in

Fig. 3). Alternatively, the filters selected by the filtering system may intercept and filter the messages before they are passed onto the characterization module 310.

The context-based information filtering system 124 can be configured to characterize and filter any information that is analyzable by a computing device. The precision of the filtering is in part determined by how much information about the information is available. At a minimum, the system can use characteristics that can be inferred from the information itself, such as message characteristics for received messages. Examples of such message characteristics include the following:

- source – a type or instance of carrier, media, channel or network path;
- destination – address used to reach the user (*e.g.*, a user typically has multiple address, phone numbers, etc.);
- content of the message (parsable or described in metadata);
- data format type;
- arrival time;
- size; and/or
- previous messages—inference based on examination of log of actions on similar messages.

In addition, many message formats explicitly include message-characterizing information, which can provide additional filtering criteria. Examples of message-characterizing information includes the following:

- title;
- originator identification (*e.g.*, email author);
- origination date & time;
- routing (*e.g.*, email often shows path through network routers);
- priority;
- sensitivity – security levels and permissions;
- encryption type;
- file format – might be indicated by file name extension;
- language – may include preferred or required font or font type;
- other recipients (*e.g.*, email cc field); and/or



- required software.

Moreover, the information source may add any additional kind of characterizing metadata to the messages, and such metadata can be used to determine message appropriateness. Examples of such metadata include the following:

- offer characteristics – a solicitation for goods or services can include an explicit description of the offer, such as specific product attributes like size, color, materials, cost, availability, warranty, external reviews, etc. with adequate detail, messages can be presented to the user corresponding with a narrowly defined user context;
- sender characteristics – their context can include arbitrary attributes (mood, location, who they are in presence of);
- certification – a trusted indication that the offer characteristics are dependable and accurate; and/or
- recommendations – outside agencies can offer opinions on what information may be appropriate to a particular type of user or situation.

In the illustrated embodiment, the characterization module 310 characterizes the user's context based on information received from a user context module 312. Information pertaining to the user's environment is captured and passed to the user context module 312 for analysis. The user context module 312 creates, maintains, and updates one or more models pertaining to the user's context that integrate multiple areas of a human's computerized environments, such as physical, mental, computing, and data. In the illustrated implementation, the user context module 312 gathers information on (1) the user's physical environment from module 314, (2) the user's mental environment from module 316, (3) the user's computing environment from module 318, and (4) the user's data environment from module 320. Various sensors 322 (*e.g.*, environment sensors 224, user sensors 222, etc.) monitor the environment parameters and provide data to the modules 314-320.

The physical environment module 314 generates information pertaining to the user's present location (*e.g.*, geographical, relative to a structure such as a building, etc.), the current time, and surrounding objects within a perceivable proximity that may

interfere or distract the user from focusing attention on the incoming messages. Other possible information related to the physical environs include people, user's body state, or anything that can be externally sensed by computer-connected transducers or inferred by examining other parts of the context model. As an example, a user with a wearable  
5 computer may be traversing through a structure that is very dark and cold and that requires full visual and auditory senses on the surrounding objects within the structure. While in this state, the user most likely does not want to receive an advertisement for a sale on sunglasses.

The mental environment module 316 generates information pertaining to  
10 the user's likely intentions, their preferences, and their level and quality of attention. Other information may include mood, topic of conversation, or any self-characterization that is meaningful to the user, provided by an observer, or inferred by examining other parts of the context model. For instance, the mental environment module 316 may use data from a pupil tracking sensor or head orientation sensor to identify a direction or object on which the user is focused. If the user appears to be focused on administrative items presented on a heads-up display, then the user context module 312 might determine that it is safe to present a message. Another example is to use preference metrics programmed by the user to differentiate among the urgency of certain incoming messages. For instance, the user might specify that all messages from his family are highly important, whereas all messages from a book retailer are marked with the lowest  
20 level of importance. Filters could then filter messages based on these defined importance levels.

The computing environment module 318 generates information pertaining to the computing capabilities of the client, including available I/O devices (*e.g.*,  
25 keyboards, buttons, microphones, cursor controllers (such as rolling, pointing, pressing, bending, turning, twisting, switching, rubbing, zipping controllers), displays, speakers, pressure devices, etc.), connectivity (*e.g.*, amount of bandwidth, security level, cost, strength, protocol, quality, media type, schedule, etc.), processing capabilities, available storage space, and anything else that the local and remote computing environment can  
30 self-characterize. The data environment 320 generates information pertaining to the data

and software resources on the client computer, including the communication resources, applications, operating system, and data.

The user context module 312 takes the information generated by the environment modules 314-320 and generates a cohesive set of user conditions based on one or more context models. The model(s) can be broad, dynamic, explicit, and capable of being understood by arbitrary software modules. Here, the set of conditions are provided to the characterization module 310 which is able to comprehend those conditions. Alternatively, the architecture could support a dictionary that may be accessed by any module wanting to understand the meaning of a specific attribute or condition produced by the models.

The characterization module 310 uses the conditions to identify a set of one or more filters to apply to the incoming messages 302. The set of filters is dynamic and can be changed as the user's conditions change. New filters are added or created, while others are dropped as the user's context evolves. For example, in some embodiments each filter could specify a set of one or more context model attributes, and include for each of the specified attributes one or more possible values for the attribute. Such a filter would be selected as matching the current modeled context (and thus used to apply to incoming messages) only if all the specified attributes have current values in the modeled context that are one of their possible values. Alternately, each specified attribute could include an indication of whether it is required, and in such embodiments the filter can still match the current context even if the current values for non-required attributes do not match the specified possible values. Filters could also specify inter-attribute logic.

As described in greater detail in U.S. Patent Application No. \_\_\_\_\_ (Attorney Docket No. 294438008US1), filed April 2, 2001 and entitled "Thematic Response To A Computer User's Context, Such As By A Wearable Personal Computer," multiple themes can also be defined for a user and used to present appropriate information and functionality. In some embodiments, filters used to analyze incoming messages or other information can each be associated with one or more defined themes, and the filters selected for current use can be one or more of the filters that are associated with one or more of the themes that are determined to match the current context.

In one implementation, the filters 126 each specify criteria to be matched against the content in the incoming messages 302. The messages are selectively allowed or blocked depending on whether the content matches the filter criteria. The messages themselves may be constructed with only the content (*e.g.*, product or service information in an advertisement), and if so the filter or characterization module parses the message for key content to match against the filter criteria. As an example, the message may be a text or radio message that touts a particular product, and the filter or characterization module could parse the message using a text parser or a speech recognition module to identify the product name, type, cost, and so forth.

Alternatively, the information source may supply metadata to aid the filtering and thus avoid the parsing of the messages. In addition, the source-provided metadata can further describe the most appropriate use of the message by characterizing the intended or preferred message recipient's by including target context attribute value ranges. Those skilled in the art will appreciate that in some embodiments the use of such source-provided information may be limited to sources and/or messages that are trusted (*e.g.*, on the basis of being certified by an independent third-party) and/or the system could take additional steps to verify the accuracy of the source-provided information.

The message contents or metadata (from the parser or the source) is then compared to the filter criteria of the currently selected filters. In this way, the filters can act as a proxy for the user's current situation and preferences in order to determine whether messages are appropriate, and thus should be blocked or kept. The characterization module 310 may further examine the context data to determine whether to save the messages for later or pass them onto the user immediately. Message filtering can thereby be accurately and generally based on context parameters that model and evaluate whether messages are appropriate under a set of circumstances.

The characterization module 310 may utilize various characterization models to implement the filter comparisons. One way to characterize messages is to assign a numeric value to a set of message attributes and then compare that numeric value to a table of minimum necessary values or combination of values in order to determine

whether a particular message is suitable under the circumstances. An example set of message attributes includes the following:

- Message Desirability/Appropriateness – This attribute examines whether message conforms to enumerated or inferred criteria in the filters. For example, a filter might present only the high-priority messages from work-related contacts when the user's context model indicates that they are not working. As another example, the user may specify types of message as desired, such as any information on a particular product or tickets to a concert. These explicitly stipulated characteristics could be used to present only solicitations for the specified product or tickets. Such criteria may also be inferred, such as whether the user's sex, height, or weight has a bearing on a piece of clothing he/she might be interested in purchasing. Also, the characterization module may indicate past preferences regarding purchase decisions, like brand of product, store from which the product was purchased before, and so forth.
- Message Need – This attribute helps determine that a message should be presented even though there is no corresponding interest characterization. For example, a message concerning the safety of a product, or a hazard to the safety and health of the user may be deemed appropriate to present.
- Urgency – This attribute assists in determining whether messages should be presented immediately, or after some delay.
- Sensitivity & Required Privacy
- User's Available Attention

Consider the following examples of how the filtering system might operate. Suppose an unsolicited advertisement for a cooling drink is received at the client. The user context module, either in response to the advertisement or as a part of separate processing, determines whether the user is thirsty based on various environment data and physiological state information. For example, the user's

temperature, the surrounding temperature and humidity, and the user's past behaviors can all be used to determine whether the user is thirsty. In addition, other context information can also be employed to determine whether the advertisement is appropriate for presentation (e.g., an indicated or learned preference that the user does not want to receive advertisements while they are exercising). The various relevant context information is passed to the characterization module 310, which then applies one or more filters suitable for controlling whether a cooling drink advertisement should be presented to the user at this current time.

As another example, the filtering system 124 may delay presentation of an acceptable message that matches the filter criteria until a filter 124 or the characterization module 310 determines that the user context matches the presentation requirements of the message. Suppose the user's current context suggests that the user is not able to give attention to incoming messages. For instance, the user is driving a car and is not able to view the message at this particular time. Now, suppose that the user receives an incoming message that is of interest to the user (e.g., a one-time sale on travel tickets to a desired destination). The filters and characterization module allow the message to successfully be received, but the user's context determines that the message cannot be displayed yet. Accordingly, the filtering system 124 accepts the message, but delays presentation until the user's context has changed (e.g., stopped driving) and the user is in a better condition to receive the advertisement. When the context finally changes or when the user specifically notifies the characterization module to present the message, the user can experience the message without interruption or distraction.

As another example, the filtering system 124 may employ a multi-tier filtering process using a hierarchy of multiple filters. The filters may be directed to different sets or categories of messages, with the collection of filters effectively finding a subset of messages acceptable to each of the filters. The characterization module can construct a filter composite as conditions warrant to effectively and dynamically produce unique filters as the user's context dictates.

Such construction of multi-tier filters could be performed in a variety of ways. For example, in addition to specifying information (e.g., context attributes) that

indicates whether a filter matches the current context, in some embodiments filters could additionally specify inter-filter information that assists in determinations of whether two or filters are compatible. Such inter-filter information could indicate a directional relationship, such as one filter providing an output indication that is used as input as another filter, or merely that two filters are not inconsistent (e.g., a first filter that prohibits advertisements about specified products would be inconsistent with a second filter that indicates a desire to receive advertisements about those same specified products).

In addition, information about categories or groups of filters, and/or of hierarchical relationships between filters, could also be used when dynamically generating multi-tier filters. For example, if a first filter selects a group of information as appropriate that is a proper subset of the group of information selected by a second filter (e.g., a second filter that selects any drink advertisement as appropriate and a first filter that selects only diet soda drink advertisements as appropriate), use of the second filter in conjunction with the first filter would be redundant – this type of inter-filter relationship may be ascertainable based on a hierarchical relationship between the filters. Similarly, in order to determine filters that are compatible, it may be appropriate in some embodiments to select only filters that are members of a common group or category, or to instead select only filters that are each from a distinct group of category. Examples and uses of hierarchies, groups and categories are discussed in greater detail in U.S. Patent Application No. \_\_\_\_\_ (Attorney Docket No. 294438008US1), filed April 2, 2001 and entitled “Thematic Response To A Computer User’s Context, Such As By A Wearable Personal Computer,” and such information can similarly be applied to filters.

Fig. 4 illustrates a filter composite 400 that is made up of four filters 402, 404, 406, and 408. As an example, the first filter 402 might specify generic characteristics of suitable messages, such as the product/service characteristics for advertisements. The second filter 404 might specify user characteristics and preferences. For instance, the user might only want to receive advertisements while running errands. The third filter 406 might describe certain user situations in which messages are appropriate. In this example, the user might only want to see or hear advertisements for

products when proximally near a store. The fourth filter 408 might pertain to the client computer capabilities. For instance, the client may be a communication device that can only receive text messages, as opposed to graphical or video messages. A message satisfying all the filter criteria of the filter composite 400 is then stored and/or presented to the user.

### **Message Filtering Process**

Fig. 5 shows an example process 500 for filtering messages according to a user's context. The process 500 may be implemented in software, or in a combination of hardware and software. As such, the operations illustrated as blocks in Fig. 5 may represent computer-executable instructions that, when executed, direct the system to display virtual information and the real world in a certain manner.

At block 502, messages (*e.g.*, unsolicited messages) 302 are received at the client 106 from the information sources 102. At block 504, the context-based information filtering system 124 selects one or more filters 126 to apply to the incoming messages. The characterization module 310 of the filtering system 124 is frequently updating the set of filter(s) 126 to use for incoming messages according to the environment conditions of the user. As the messages are received, the current set of filter(s) 126 is applied to the messages (block 506).

At block 508, the characterization module 310 determines whether the messages match the filter criteria. If the messages fail to match the criteria (*i.e.*, the "No" branch from block 508), the messages are discarded (block 510), or could instead be stored for later consideration (*e.g.*, explicit user review or more detailed automated processing). Alternatively, if the messages match the criteria (*i.e.*, the "Yes" branch from block 508), the characterization module 310 evaluates the user's current context to decide whether the user should currently experience the acceptable messages (block 512).

At block 514, the characterization module 310 determines, based on this evaluation of the user's context, whether to currently present the message. If the user context indicates that it is appropriate to present the message to the user (*e.g.*, the user an agreeable state to review the messages), the characterization module 310 forwards the



message(s) for immediate presentation (block 516). Presentation may be in the form of displaying a text message, rendering a graphical or video message, playing an audio message, releasing an aroma of a olfactory-enabled message, and the like.

If the user context indicates that it is not appropriate to currently present the message (*i.e.*, the “No” branch from block 514), the characterization module 310 stores the message(s) for later presentation (block 518). When messages are stored, the characterization module 310 can continue to monitor the user context for any changes that might allow presentation of the stored messages. The characterization module can also in some embodiments provide notification to the user to inform them that one or more messages are available for presentation (*e.g.*, an indicator shown on the display or a subtle background sound that is played), thus allowing the user to explicitly select the messages for presentation when the user desires.

There are a number of ways that the content may be presented to the user. Content may be displayed, audibly played, released as an odor, output as a tactile feedback, and the like. Additionally, the choice of a suitable presentation mechanism may be determined based on the user’s context. That is, after the content passes the filters and is determined to be presentable at this moment to the user, the characterization module 310 might still further decide how best to present the information. For example, a user who is preoccupied watching objects or their surroundings may not wish to see any messages, but might be inclined to listen to them. Thus, the characterization module might determine to play any audio messages that might be ready for presentation. In other situations, the user might be more inclined to see the information on a display or watch a monochrome.

### **Message Distribution System with Intermediary**

Fig. 6 illustrates an exemplary message distribution system 600 in which an intermediary entity 602 acts as a distributor of messages on behalf of an information source. This system 600 is particularly suited for the commercial advertising market in which the information sources consist of advertisers 604 that are attempting to locate users who might be interested in their products and services. The system 600 supports a

revenue model for deriving revenue based on the service provided by the entity 602. The revenue model includes a variety of possible fees for different levels of service.

The entity 602 performs services related to matching messages from information sources with suitable users who might be interested in the messages. The entity 602 maintains a registry 610 of interest-rules that are collected and maintained on behalf of individual users. The rules are correlated with the users to represent, for a given user, a composite of the user's desires, preferences, past behaviors, presentation surfaces capabilities, user demographics, and so forth. The entity 602 also maintains a message store 612 to store messages (*e.g.*, advertisements) received from the advertisers. A match module 614 compares the messages in the message store 612 to the interest rules in registry 610 to identify users who might be interested in selected messages.

Each client 620 is equipped with an interest module 622 that generates rules indicating which items may be of interest. The interest module 622 creates these rules based on specific user-defined preferences 624 and traits/characteristics gathered by the characterization module 310. The client 620 occasionally submits the rules for that user to the entity 602 for storage in the registry 610.

When advertisers look for interested consumers, they can send their advertisement (optionally with metadata) to the entity 602. The advertiser may be asked to pay a subscription or handling fee to the entity for access to likely consumers, and in some embodiments the intermediary entity provides a payment module to facilitate the exchange of payment between the intermediary, the clients, and the information sources. The entity 602 stores the received advertisement in the message store 612, and then invokes the match module 614 to compare the advertisement with the rules in the interest-rule registry 610. If there is a match, this indicates that a particular user might be interested in the advertisement. There may be a fee associated with matching the message to various users based on their interest rules.

When a message is determined to be potentially of interest to the user, the entity 602 delivers the message to the user's client 620. The client-side characterization module (not shown in Fig. 6) then determines, based on the user context, whether to present the message or store it for later presentation. In some embodiments, the filter

mechanism could rely on the intermediary to provide information that is appropriate to the current context, while in other embodiments the filter mechanism could perform additional processing using current filters to determine whether to accept delivered messages. If the message is presented to the user, the advertiser may be asked to  
5 compensate the user and/or the entity 602. If the advertiser sells something to the user, the advertiser may be asked to provide an additional fee to the entity 602 for locating a buyer.

The entity 602 may be configured to perform other services that are beneficial to the users. The entity 602 may act, for example, as a trusted agent that  
10 examines messages submitted from the information sources to determine whether the messages adhere to acceptable standards of message characterization. The entity could be responsible for evaluating message format, structure, content, metadata, and the like to determine conformance with desired standards. As a trusted agent, the entity 602 can certify messages as being in compliance with these standards for the benefit of the user filters. Additionally, the entity could certify and verify information sources, to facilitate swift examination of messages from certified sources. Certified messages or messages from certified sources may be issued a secure certificate to indicate to the user filters that the message metadata can be trusted to accurately characterize the message, and to ensure that the filter will recognize the certificate's meaning. The message filter can respond to  
20 the secure certificate by automatically assigning a higher confidence rating to the message.

The interest rules for the user can be generated in a variety of ways. One example technique for generating such interest rules is described in greater detail in U.S. Patent Application Serial No. \_\_\_\_\_ (Attorney Docket No. TG1-004US), filed April  
25 2, 2001 and entitled "Soliciting Information Based On A Computer User's Context." In addition, other techniques for generating themes and theme-related information are described in greater detail in U.S. Patent Application No. \_\_\_\_\_ (Attorney Docket No. 294438008US1), filed April 2, 2001 and entitled "Thematic Response To A Computer User's Context, Such As By A Wearable Personal Computer." Similar  
30 techniques to those disclosed for generating user interest rules and theme-related

information, either in combination or separately, can also be used to allow users to create and modify filters for use in analyzing received information.

In addition, in the same manner that advertisers or other information sources can provide messages or other information that may be of interest to users, such sources could also supply filters to be distributed to users via the intermediary entity or instead directly to users. In some situations such filters could be distributed as part of a business model, such as by selling the filters to users. In other situations the filters could be distributed for the convenience or benefit of the information sources (*e.g.*, a product/service provider or related advertiser providing a filter that selects as appropriate only those advertisements that are for those products or services), and in some embodiments the information sources could compensate the users for accepting and/or using the provided filters. Such filters could also be provided to users based on a group to which the user belongs (*e.g.*, filters distributed by an employer to all employees or to employees meeting specified criteria) and/or based on a geographic location of the user. In some embodiments, filters are associated with themes (and are thus one type of theme-related information for those themes), and can be distributed in conjunction with the distribution of those themes.

In addition to (or instead of) automated selection of filters, in some embodiments users can explicitly indicate appropriate filters. For example, in response to receiving information that was determined by automatically selected filters to be appropriate, the user could indicate that the information was instead inappropriate (or was presented in an inappropriate manner). The user could also indicate that one or more of the automatically selected filters should not have been selected and/or that another filter that was not selected should have been selected. In this way the user can explicitly adjust the information that is determined to be appropriate.

Moreover, in embodiments in which the user provides explicit indications of appropriateness (whether by selecting filters or by indicating that filters and/or presented information is appropriate or inappropriate), the system can use automated learning techniques to enhance future selections and use of filters. In this manner, a filtering system can be continually customized to a particular user, thus enhancing the

usefulness of the filtering. Additional details about the use of such automated learning techniques are provided in U.S. Patent Application No. \_\_\_\_\_ (Attorney Docket No. 294438026US), filed April 2, 2001 and entitled "Improving Contextual Responses Based On Automated Learning Techniques", which is a continuation-in-part of U.S. Patent Application No. 09/216,193, of U.S. Patent Application No. 09/464,659, and of U.S. Patent Application No. 09/724,902, and which claims the benefit of provisional U.S. Patent Application No. 60/193,999 and of provisional U.S. Patent Application No. 60/194,123. Each of these applications is hereby incorporated by reference in their entirety.

In addition to using modeled context information to select filters or dynamically generate filtering capabilities, such context information can more generally be used in some embodiments to control the sharing of a user's context or other computer-controlled information with other users and/or computing devices. Such sharing of information can be conceptualized as the filtering of outgoing information (rather than incoming information), and arbitrary groups of context attributes can be specified as constraints on the sharing. For example, based on the user's modeled location, physical activity, presence of particular other people or objects nearby, and emotional state, it may be determined that certain groups of user context information will be made available to people in the user's family, that other groups of user context information and other stored information will be made available to co-workers, and that no one other than the user may have access to yet other groups of information. In particular, in some embodiments information sharing filters will be defined and used in a similar manner to that described above for received information filters. Alternately, existing mechanisms for grouping or using groups of context attributes can be extended to include information sharing filter information, such as defined themes or received information filters.

There can be multiple tiers of information sharing (or "permission") characterization, and such tiers can be selected in various ways. One example of tiers of information include information that is available to the user themselves, to other users in one or more groups or categories to which the user belongs, to users in other groups or

categories to which the user does not belong but has a defined relationship, and to everyone. Those skilled in the art will appreciate that these tiers can be expanded or contracted in various situations – for example, as noted above, a user may belong to many groups (*e.g.*, family, club, work, etc.) that each have different permission schemes.

5 Similarly, in a specific group of applications such as multi-player games, the other groups can include competitors, whether individuals or teams, and additional tiers could be added, such as for non-player spectators or for game creators/administrators. In a business setting, other groups may include partners, suppliers, customers, etc. Such permission tiers can also be used when receiving information from others that are  
10 employing such information sharing filters, such as to receive team information from other members of your team in a multi-player game setting, whether the information is aggregated (*e.g.*, by a third party, such as a team leader) or specific to each of the other members.

Permission information can also be specified in various ways. Since a context attribute can be shared as a data item, permission information related to that context attribute can be specified with an explicit characterization of permission and/or an associated permission algorithm (*e.g.*, as an extended property of that context attribute), or the value of the context attribute can instead be treated as any other information that potentially may be distributed. In particular, any such information can  
20 be characterized with one or more properties (*e.g.*, based on the information content or other related information) as to who/what/when/where/how/why that information can be shared. For example, it may be desirable to have all context information and other information that is related to my current and historic physical health be characterized as being available to medical professionals providing assistance during a medical  
25 emergency. In addition, in some embodiments a user may not have authority to examine and/or modify the permission scheme or information associated with their context attributes or other information. An example of this may be in an institutional setting where the system is provided for a purpose that is controlled by the institution.

In addition to the general applicability of the disclosed techniques for  
30 selecting appropriate information, various specific uses of the previously disclosed

techniques and other related techniques follows. For example, computer gaming environments provide situations where modeled user context information can be used to significantly enhance the gaming experience, such as by selecting appropriate game-related information or functionality based on the user's context and presenting it in an appropriate manner. When some or all of the game participants use wearable computers, even greater enhancements are available, as discussed in greater detail below.

Thus, in some embodiments techniques for determining various details of the user's current context (*e.g.*, the physical location of the user in the real world and the user's current mood) and interacting with the user in various ways based on that context information. In some embodiments, such techniques could be provided directly by gaming applications designed to implement the techniques, while in other embodiments the techniques could be added to existing games by an external component (*e.g.*, provided by a third-party).

One way that gaming applications can be enhanced is by augmenting the user's reality (*i.e.*, integrating computer-provided information into a user's perception of their physical reality) via the I/O devices of the user's computer. The enhanced I/O device functionality that is typically available on wearable computing devices is one reason that such wearable devices can in some situations provide enhanced functionality not available to standard computing devices that lack the I/O device functionality.

The ability to augment reality in this manner provides various benefits in both gaming and non-gaming applications. For example, in gaming applications the game-related information (and thus the corresponding gameplay) can be layered on top of real life, and the game player's perceptions can be enhanced by the wearable computer. In a similar manner, the real world can be brought into the game in various ways. For example, virtual objects can be made visible to the player (objects, scenes, matte painting, etc.) but not to other nearby non-players. In addition, avatars (whether real, computer-generated players, or opponents) can similarly be presented, and such avatars can additionally react to the user's context (*e.g.*, presentation of an avatar can be based on the player's physiological state, such as displaying an avatar in a different manner if the user is startled or afraid). An avatar representing a player can also be modified based on the

context of the player, such as making the avatar more visible if the person is scared or if the person or avatar is moving too fast.

Other examples of enhanced gaming (or non-gaming) functionality using augmented reality include displaying virtual information, such as virtual footprints that reveal the passage of a person through an area so as to allow tracking/trailing someone in the game or in real life or to allow the user to retrace their own path. In a similar manner, virtual graffiti tags (*i.e.*, computer-generated markings visible only to participants in the gaming system) can be displayed, providing a variety of possible games and uses (*e.g.*, playing ‘Tag’, or playing a virtual ‘Go’ game played in any agreed boundary, such as by using city blocks as a playing game and seeing the Go pieces from an above-city view, etc.)

While there are existing examples of presenting a visible logging of where someone or thing has been on a representation of terrain, only very limited information is typically conveyed. By providing a system where arbitrary information can be associated with the logging, the user can enjoy the benefit of richer and more appropriate information that is presented in an optimal manner. For example, the display of the virtual footprints can be overlaid on the user’s vision of the real world in the same location as the actual steps were taken. Other functionality can also be provided, such as varying footprint visibility based on the user’s context (*e.g.*, the skill of the prey and/or tracker, or having footprints that fade over time). Alternately, an entertainment or training application can use a log of a route of a well-informed person as a guide to others, and information such as the guide’s comments or recommendations to supplemental information can also be presented (*e.g.*, as iconographic controls within the route, or as audio information when the location is achieved). Individuals and institutions can use individual or aggregates of individual paths to determine used routes, which can assist in identifying areas of interest and congestion. By allowing such a log of locations and routes to be public, a “Public Scrapbook” can be created over time. In addition, by being able to see a route overlaid on the view of the physical world, a user can easily determine if they have been somewhere before (*e.g.*, to avoid being lost or to find a way through complex terrain).



The ability of the computing devices of the various players to communicate also provides enhanced functionality. For example, distributed observations can be provided by a system of connected players, such as by augmenting a player's senses with information from teammates or allowing a player to hack into other team's members and experience their environments. Remote tele-presence sensing can also be provided, such as providing a view to a user of a lookout in a tower. In addition, dynamic aggregation of players can be used to determine if there are enough people available to play a game at a given moment, either locally or worldwide depending on the nature of the game. The use of context information about such other players allows for the choosing of opponents and team members based on the capabilities, interests, or other arbitrary characteristics of the players. In addition, game rules (*e.g.*, defaults) can be modified based on communications between players (or by selecting a group of players with common interests).

Thus, as noted above, augmented reality allows games to react to real world events, such as by determining player positions in the game via the physical location of the user (*e.g.*, via available locating devices such as GPS, EGPS, etc.). More generally, the context attributes that are used to characterize a user's virtual world could be based on the context attributes from the physical world – thus, a computer-generated person or object shown in an augmented reality context could have the same characteristics as a real person. One application of such attribute mapping would be the control of a digital representation of a user by basing its appearance in part on the user's context model, thus allowing the digital representation to respond to changes in the user's physical world. For example, in response to a user becoming fatigued in the real world, the corresponding avatar's strength or speed could similarly be reduced, or an avatar's posture could change along with that of the user. A context attribute such as user fatigue can be measured in various ways (*e.g.*, the speed of user movement, stooping, speech slurring, etc.) In addition to basing an avatar's appearance and capabilities on those of its corresponding user, other game-provided functionality can also be altered in response to changes in a user's context. For example, the capabilities provided to a user could be varied based on the skill level of the user. In addition, games can be dynamically optimized during play

based on the changing user context (*e.g.*, changing skill level or tendencies of the player that are learned by the game), whether automatically by the game or by a game organizer.

Games can also react to or incorporate the real world in other ways. For example, users' movements can be incorporated fully into the game, and are not limited to the constraints of a specialized input device such as a Dataglove or a device that senses hand movements immediately proximate to the device (*e.g.*, via analysis of video image from a camera that can see a user's arm movements, or via motion sensing devices on various limbs of the user). Tracking and using such movements allows fighting games to include more than just hand-eye coordination, and can allow simulated use and movement of virtual weapons (*e.g.*, guns and sabers) or virtual body parts (*e.g.*, playing karate against the computer).

In a similar manner, augmented reality allows the gaming experience to impact the real world in various ways. For example, if a user is wearing a goggle display device or a computer-controlled brace on a limb, the player's vision could be occluded or the use of the limb restricted based on events that occur in the game. In addition, using output devices such as goggle display devices and earpiece speakers/noise-canceling devices, the presence or appearance of a player can be altered for other players. Moreover, new senses can be generated and supplied to the player, and existing senses can be enhanced or remapped. For example, a "Hide & Seek" game is more challenging if the players have infrared or night-vision capabilities, and sonar allows a user to see behind themselves. Similarly, augmented hearing or sight can provide new abilities to a user, or information (*e.g.*, a ringing sound) could be generated to represent an entirely new sense (*e.g.*, the ability to sense the nearby presence of Evil).

Such enhancement capabilities can benefit a wide variety of currently available games, and can additionally make possible new games. For example, an "Assassin" game could be played by any group of people with wearable computing devices at any time (*e.g.*, strangers at an airport dynamically identify each other as interested in playing and then begin playing using virtual weapons). As noted above, player appearance can be altered with appropriate output devices, and additional information can be provided via linked gaming computing devices (*e.g.*, an Assassin is

within 100 yards). “Team Fortress” could similarly be played in the real world, allowing real-time collaboration between players, as could “Virtual Tennis” between two players anywhere in the world. Similarly, “Rally”, a scavenger hunt using player context (e.g., current location) to generate clues, and virtual gambling that is constantly available  
5 (whether individual or group games) are other examples.

In addition, new capabilities could be provided, such as an ability to plant virtual sensors that detect other game players or to plant virtual landmines that reveal players when triggered (e.g., visually augmenting the revealed player’s presentation and/or providing location and other context information about the player to other  
10 players). In addition, sensors could be provided in physical devices, such as a Bluetooth-enabled golf ball that assists users in recovering lost balls. Moreover, if various mental user context information can be detected (e.g., via choices signaled by Hz on-screen-frequency-selection and/or via biofeedback), games could react to such information. In addition, a variety of types of interesting information (both context attribute information and other information) could be provided to external audiences (or to team members or competitors) in a similar manner to cameras in NASCAR autos. This allows observation of games that are currently unobservable, such as “Half-Life” or various other games viewed with a wearable computing device. By providing relevant information to a game organized/manager, that person can opportunistically manage on-going games.

Figure 9 is a block diagram illustrating an architecture for enhancing gaming environments by providing appropriate information and functionality based on a user’s context. As is shown, various components execute in memory of a user’s gaming computing device, including one or more game applications 938. Various optional game enhancement components 941-949 are also illustrated, although those skilled in the art  
25 will appreciate that in alternate embodiments only some of the components may be present. In addition, in other embodiments the game applications themselves will incorporate the desired enhancement functionalities. While the illustrated embodiment also shows a received information filtering component 934, information sharing component 936, and context modeler component 932 executing in memory, in other  
30 embodiments some or all of these components may not be present. The information

filtering and sharing components, as well as various of the game enhancement components, can communicate with various other computing devices as illustrated.

Figure 10 is a flow diagram of an embodiment of a Gaming routine. The routine provides various enhanced gaming functionality by using context information of the user, as is illustrated and is described elsewhere.

As noted above, the use of various context information can provide a variety of types of game enhancements. Examples of relevant types of context information include the following:

#### Physical Characteristics

- Fatigue – if the player is tired, this can be used to change the capabilities of their game character by making it slower.
- Dexterity – someone with limited physical skills may be programmatically associated with characters of reduced skill, thereby giving a realistic expression to their capability. Conversely, their input device responses may instead be enhanced, thus giving them opportunities to compete on a more equal and satisfying basis.
- Location – where a player currently is, or will be. This can be used for games that involve the real world, such as geo-caching, or games that require access to physical facilities. This is also useful for games and competitions that are races through physical locations.
- Availability – when a player can play, for how long. Also, how much they are willing to pay or risk.
- Payment – indication that player or group has paid entrance fee or dues.
- Insurance – indication that player or group has necessary insurance.

#### Mental Characteristics

- Skill – a player's or teams' skill (e.g., based on self-characterization, automated handicapping system, 3<sup>rd</sup> party rating, etc.) can be used to match them with appropriate adversaries. This can also be used by a spectator to choose what matches to watch.

- Memberships – enumeration of which organizations or groups a player belongs to. This can include certifications by gaming, safety or other organizations.
- Familiarity – in contrast to Skill, which relates to proficiency, Familiarity provides an indication of a potential player's ability to play the game. For example, are they familiar with the rules of the game?
- Interest – a player can indicate, or have stored an indication of, the characteristics of games that they are interested in joining, observing, officiating, reviewing, buying, or selling.
- Current & Past Scores – a log of an individual's or teams' performance can be maintained and shared.
- Age – useful in determining appropriateness of game content. Since game environments can dynamically respond to changes in aggregate player characteristics, this attribute can be used to determine when it is appropriate to change the content from child to adult. For example, later in the evenings, after children stop playing, more adult language and themes can be employed. Alternately, when a family is simultaneously playing the same game, provide a more limited or different set of content to the children than is provided to the parents.
- Roles – Role playing games often have game characters that have persistence over multiple game sessions. Attributes characterizing the game roles, including character names, experience, and skills, can be shared.
- Language – which language(s) the player speaks
- Risk – indication of how much risk, physical or monetary, the player is willing to experience.

#### Platform Capability

- I/O Devices – to what devices does the player have access, and what are their capabilities? For example, can the player use a head-mounted

display that supports the superimposition of color computer-generated images over their view of the real world?

- Processing – how fast is the player’s computer (CPU, RAM, video cache, etc.)
- Game Capability – an enumeration of games, or levels of games, can be explicitly enumerated
- Storage – how much mass storage is available for games
- Software – what software is available

#### Data Capability

- Connectivity – what is the bandwidth and speed (sometimes called “ping time”) of the computer platform’s network connections
- Security – what protocols are supported

Various group characteristic information can also be provided and used, such as attribute permission schemes that indicate the scope of information to be shared. Attributes can also characterize or identify the different groups of users to which the permission scopes will apply. For instances, different group scopes can include the following: the player (they might be able to see all of their attributes, or instead some may be hidden or provide false information during some games), the player’s team (they made be able to see some/all of their teammates, and aggregate team attributes), the game’s organizers and officials (they could see all players attributes, for example), the game’s commentators (they could see some/all team and player attributes), and the game’s audience (they could see some/all team and player attributes).

The disclosed techniques for selecting and presenting appropriate information can be used in some embodiments to modify the normal operation of the user’s senses in various ways. Moreover, there are many examples of sensory enhancements that can benefit from a general model of user context. For example, context information can indicate when and how to map electromagnetic spectrum frequencies to human perceivable frequencies (UV, IR), sound, pressure, etc. Each spectrum frequency could additionally be mapped in a variety of ways. For example,

with respect to sound, the mapping could provide example functionalities as indicated in the following list:

- greater frequency sensitivity
- greater volume sensitivity and/or volume normalization (e.g., increase low and suppress high frequencies, or increase the compressed dynamic range such that the average sound is greater even though the peak volumes do not change)
- selective frequency suppression (e.g., low rumbles while in engine room)
- noise cancellation (not just fixed, but user-selective in type and degree)
- mapping of senses (e.g., proximity to objects is translated into sounds, a VU meter, determining the tempo of the music, etc.)

Moreover, a computing device could provide additional sound-related functionalities, such as triggering notifications or functionality when certain conditions are met (e.g., when siren occurs the audio is lowered, a SID monitor, detect when engine sound deviates from normal, a dangerous volume monitor, etc.). Sounds could also be automatically identified, such as by identifying a human speaker based on voice, a fire truck or police car, friendly versus enemy fire, etc. Sounds can additionally be characterized in various ways, such as to determine whether the gun fire is aimed in my direction, whether the fire truck is coming toward me, whether that bird call was from a real bird, how far away was that twig snap, etc.

Augmenting reality (e.g., with computer-generated images, sounds, smells, sensations, etc.) is not limited to adding or transforming indicated, measured or inferred data – instead, it can also include the blocking of real-world stimuli. For example, when light of a potentially hazardous amplitude or frequency is detected, the portion of a transparent display through which the radiation is passing can be made opaque, thereby protecting the user's eyes. Similarly, audio noise canceling algorithms can be used to protect hearing. The system can also obscure undesirable options, such as sub-optimal routes when driving at high speed in an unconstrained terrain (like desert). For workers who are be visually distracted by their environment (e.g., by people in adjoining

cubicles), non-work related stimuli could be obscured. In addition, computer controls can be hidden based on the context as an alternate way to give prominence to something in the real world (*i.e.*, by obscuring less relevant information). For training purposes, the system could eliminate incorrect choices over time to assist a user in identifying the correct choices and judging their mastery of the task. For entertainment purposes, a temporary or permanent loss of perception can be provided during a game (*e.g.*, when a player has been successfully attacked). Moreover, it is not just senses that can be modified, but also physical capabilities or cognitive ability (*e.g.*, there can be less game or context information or computing capabilities made available in order to simulate the loss of personal or equipment function).

Figure 11 is a flow diagram of an embodiment of an Augmented Reality routine. The routine augments real world information available to a user with virtual information based on the user's context, as is illustrated and is described elsewhere.

As described above, augmented reality involves presenting virtual information in conjunction with real world information that is available to the user. However, existing technology is often unable to dynamically integrate virtual images with the real world without hardware clues. Maintaining the overlaying, or visual mapping, of virtual objects on top of real-world objects is known as registration. One example of registration is superimposing a virtual image of a wire over the real-world view of the wire, and having that virtual wire remain properly associated with the real wire while the user's head may be moving. Technology can perform this by adding visual clues to the hardware, such as dots on the circuit board containing the wire. Another example of the difficulty in managing the overlay of virtual info onto the real world is the difficulty formatting computer controls without conflict with viewing the real world.

While registration functionality is difficult to provide using existing technology, the use of current context information about the user can greatly facilitate the providing of registration without the use of hardware cues. Thus, some embodiments of the invention include a context-aware system that models the user, and uses this model to present virtual information on a display in a way that it corresponds to and enhances the user's view of the real world.



In one embodiment, the system displays information to the user in visual layers. One example of this is a constellation layer that displays constellations in the sky based on the portion of the real-world sky that the user is viewing. As the user's view of the night sky changes, the system shifts the displayed virtual constellation information with the visible stars. This embodiment is also able to calculate and display the constellation layer during the day based on the user's location and view of the sky. This constellation information can be organized in a virtual layer that provides the user ease-of-use controls, including the ability to activate or deactivate the display of constellation information as a layer of information.

In a further embodiment, the system groups various categories of computer-presented information related to the commonality of the information. In some embodiments, the user chooses the groups. These groups are presented to the user as visual layers, and these layers of grouped information can be visually controlled (*e.g.*, turned off, or visually enhanced, reduced) by controlling the transparency of the layer.

Another embodiment presents information about nearby objects to the user in a manner that is synchronized with the real world surroundings. This information can be displayed in a variety of ways using this layering invention of mapping virtual information with the real-world view. One example involves enhancing the display of ATMs to a user, such as via a layer showing streets and ATM locations or a layer that displays ATMs near the user. Once the user has found an ATM, the system could turn off the layer automatically, or instead (*e.g.*, based on the user's configuration of the behavior) could simply allow the user to turn off the layer.

Another embodiment displays a layer of information, on top of the real-world view, that shows information representing the path the user traveled between different points of interest. Possible visual clues (or "bread crumbs") could be any kind of visual image, like a dashed line, dots or footprints, to represent the route the user traveled. One example involves a user searching a parking garage for a lost car. If the user cannot remember where the car is parked and is searching the parking garage, the system can trace the search-route and help the user avoid searching the same locations by displaying that route. In a related situation, if the bread-crumbs trail was activated when

the user parked the car, the user could turn on that layer of information and follow the virtual trail as it displays the information to the user in real-time, thus leading the user directly back to the parked vehicle. This information could also be displayed as a bird's-eye view, such as by showing the path of the user relative to a map of the garage.

5 Another embodiment displays route information as a bird's-eye view by showing a path relative to a map. This information is presented in overlaid, transparent, layers of information, and can include streets, hotels and other similar information related to a trip.

10 The labeling and selection of a particular layer can be provided to the user in a variety of methods. One example provides labeled tabs, such as those on hanging folders, that can be selected by the user.

The system can accomplish the presenting of virtual information on top of real-world information by various means. Three main embodiments include tracking head positions, tracking eye positions, and real world pattern recognition. The system can also use a combination of these aspects to obtain sufficient information.

20 The head positions can be tracked by a variety of means. Three of these are inertial sensors mounted on the user's head, strain gauges, and environmental tracking of the person. Inertial sensors worn by the user can provide information to the system and help it determine the real-world view of the user. An example of inertial sensors includes jewelry that detects the turns of a user's head. Strain gauges (*e.g.*, embedded in a hat or the neck of clothing) measure two axes: left and right, and up and down. The environment can also provide information to the system regarding the user's head and focus, such as via pattern-matching information of the user's head. This can occur from a camera watching head movements, like in a kiosk or other such booth, or any camera that  
25 can provide information about the user. Environmental sensors can also perform triangulation based on one or more beacons transmitting information about the user's head and eyes. The sensors of a room or a car can similarly triangulate information about the user and present that information to the system for further use in determining the user's view of the real-world. In addition, the reverse mechanism also works in which  
30 the environment broadcasts information about location (or distance from one the sensors

in the environment), such that the system in that embodiment of the invention can perform the calculations without needing to broadcast information about the user.

The user's system can also track the eye positions of the user for use in determining the user's view of the real world, which can be used by the system to integrate the presentation of virtual information with the user's view of the real world.

Another embodiment involves the system performing pattern recognition of the real world. In this embodiment, the system's software dynamically detects the user's view of the real world and incorporates that information when the system determines where to display the virtual objects such that they remain integrated while the user moves about the real world.

Existing augmented reality systems are typically designed as tools to show to the user, in an enhanced way, spatial relationships between the virtual computer-generated world and the real world. They are not designed as new means to experience the world where the computer is simultaneously experiencing the world with the user and supplementing the user's experience with information the computer holds or determines and manifesting behavior in a fashion that is integrated with the user's experience of reality. Thus, what is missing in such existing systems is that they do not incorporate an extensible context model of the user. Utilizing a context model allows more complete integration – basing augmented reality on context awareness allows the interaction to become more appropriate and the content dynamically adapts appropriately.

Thus, using context information as part of an augmented reality system allows the integration of virtual information with real world information in such a manner that the computer does not merely activate output devices that can be directly perceived by the user, but instead the computer-generated information is shown in meaningful relationship to the perception of real-world information. For example, for a computer-generated name tag to be augmented reality in the sense it is used in this disclosure, it would need to appear to the user to be shown associated with the appropriate person.

Therefore, some embodiments include some or all of the following components in order to provide augmented reality that is enhanced with context information.

- An explicit, extensible, dynamic, scalable, efficient software model of the relationship between a user's body and an arbitrarily complex representation of the user's immediate vicinity.
- Abstraction of the context model from sensors – thus, allowing an arbitrary and dynamic set of data to be used in the context model without destabilizing the system
- Determination of what content is appropriate given their real world circumstances and needs
- Logically integrating the virtual and real content from the user's point of view.
- Determination of how to best present the virtual information

Several determinations are made as part of the disclosed augmented reality system. First, what computer generated information is appropriate given the user's current context? For example, showing an image of an arrow toward North, or showing directions to the Doctor's office. Second, what is the user's current context? For example, the user may be currently heading and looking South. Alternately, the user is in a car with another passenger and can take the carpool lane, the appointment is 2 miles away but isn't for 20 minutes; and there are errands at the drugstore that can be accomplished along the way. Third, how can the appropriate information be presented so that it is meaningfully integrated into the user's perception of their world? For example, a small arrow can be displayed at the bottom of the screen pointing down. Alternately, directions can be provided that include use of carpool lanes and times that are reflective of speedier passage, with the amended directions and route drawn onto the road and the drugstore indicated as an interim destination.

Figure 7 is a block diagram illustrating an architecture for augmenting reality by presenting appropriate information based on a user's context. Various information is illustrated, as is described in greater detail below.

## User 770

Can include any person or animal that can perceive stimulus from both the real world and computer-generated.

The user maintains a bi-directional dialog with the computing environment via input and output devices, such as via gestures and utterances that are intended to convey intention to the computer, and generates context data that can include not only the gestures and utterances but also measurements, observations, narrations, third-party commentary, etc.

## Local I/O Devices 760

Can include any directly perceived output devices – an example is any video display that creates photons that are received by human eyes.

Can include any directly manipulated input devices – an example is a keyboard or camera w/gesture recognition software.

## Computing Environment 700

Can include any hardware platform or platforms, local or remote.

Provides general computer capabilities; data functions (generation, reception, processing, storage), communications, etc.

## Presentation Logic 710

Determines what content is appropriate for the user's current and anticipated contexts based on current or anticipated physical relationship between the user's body and their environment.

Further can include logic that determines how to integrate from a user's first person point of view and how to communicate the integrated virtual information to the output devices for presentation.

While the Presentation Logic component can be implemented in a wide variety of ways so that it provides the necessary data structures and logic, a simple architecture is described with respect to Figure 8. Figure 8 is a block diagram illustrating additional details of the architecture for augmenting reality, with the focus being on primary data flow rather than control information. The Presentation Logic component illustrated in Figure 8 includes several sub-components, including Format Logic 811

(optionally determines the appearance of the computer-generated data), Integration Logic 812 (determines how to meaningfully combine/integrate the content with the user's point of view of the perceived world), Content Logic 813 (determines what virtual content (e.g., computer-generation information) is appropriate to present to the user given their  
5 current circumstances, such as based on a specific task), and Environment Logic 814 (determines the spatial relationship of objects in the environment, including the user's body).

#### Format Logic 811

Responsible for presentation of virtual content, and can be implemented in  
10 various ways. An example of logic used in this step is as follows:

IF DisplayScheme = UnobstructedView

IF NOT Hazard

IF AmbientLight = HIGH

IF NewMessageOfInterest = TRUE

THEN BlinkNewMessagePeripericon

A peripericons in the illustrated embodiments is an iconic visual control that is only displayed in the periphery of the user's position. They do not distract the user from their task, while conveying limited information. They typically do not include text, as users generally do not have sufficient visual acuity to distinguish sufficient detail in the periphery of their vision.  
20

#### Integration Logic 812

Adapts augmented reality presentation based on context information, and can optionally retrieve appropriate context information so that it is available. Considers the purpose of the potential computer-generated content, and the circumstances of the  
25 user. It can also take into account hierarchical and overlapping priorities, such as from multiple themes that each match the current context. An example of logic used in this step includes:

IF CurrentVideoDisplayDevice = VideoDisplayTypeHeadStabilized

IF UserHeadPosition moves

THEN

```

IF AnticipatedUserHeadPosition = UserHeadPosition
    THEN Use AnticipatedVirtualImagePosition
    ELSE Use RecalculateVirtualImagePosition

```

5

### Content Logic 813

Provides various Virtual Objects for presentation that can include image content (text, controls, pictures, avatars, objects, landscapes, personality, etc.) and appearance. In some embodiments this component is not used if the desired computer-generated data is presented directly to the Integration Logic 812. This is typically not possible for an arbitrary number of tasks, however. Even if tasks are coordinating between themselves on which data is most appropriate to display, a general shared mechanism for them to communicate is needed and the data must take into account the user's context, which requires communications with the user's context model. An example of logic used in this step includes:

```

IF PossibleUserPostion = ProximityToHazard
    THEN DisplayHazardIndication, TellGameAboutHazard
IF CurrentUserMotion or AnticipatedUserMotion = ProximityToHazard
    THEN DisplayProminantHazardIndication, TellGameAboutHazard
IF CurrentUserMotion = CloseProximityToHazard
    THEN DisplayOnlyProminantHazardIndication, TellGameAboutHazard
IF Hazard AND Emergency AND HighPriorityMessages = FALSE
    IF Theme = Work
        IF ShowIDTags is TRUE
            IF PersonRecognizedOrIndicated = VISIBLE
                THEN DisplayNameTag

```

25

In some embodiments, the virtual objects are modeled using context information as if they were objects in the physical world.

### Environment Logic 814

Provides spatial relationships of user and environment objects, and physical characteristics of user and environment objects (*e.g.*, location, mass, motion, texture, appearance, sound, smell, temperature, etc.). Can be implemented in various ways, such as using a datastore of context attributes (whether user's private datastore or a shared context model). Can provide logic via rule/action pairs maintained in the module or provided by as a service by a separate module, such as by a remote server service.

### Task Logic 720

Can include software programs that generate user choices with feedback related to a purpose or goal. A simple implementation could be pairs of two values; Rule and Action. Rules could be potential user locations, and Actions would be what the computer should do if the location is reached. This component can interact with and make use of any software application program that programmatically communicates with a publicly known interface commands, such as Microsoft Office, Valve HalfLife, and various Internet browsers.

### User Context Model 730

In the illustrated embodiment, an explicit, extensible, dynamic software model of a user's physical, mental, data, and computer environment.

### Coordinating Tasks 740

Optionally, the system can coordinate logic with other tasks. For example, most software applications are capable of initiating and receiving instructions with other software programs, including different levels of operating systems.

### Shared Context 750

Optionally, the system can coordinate logic with other context models. Models may be of users, groups of users, or other types of context models (*e.g.*, object, location, process, etc.). Most software applications are capable of initiating and receiving instructions with other software programs, including different levels of operating systems.

The disclosed augmented reality techniques can be used in a variety of scenarios, including the following non-exhaustive list: navigation/routing, conference



calls, vehicle operation, role playing games, tours, theme parks, hazard warnings, etc. Examples of uses includes the following.

- Gaming: Multi-user applications are enhanced by the ability of gaming systems to receive abstracted descriptions from multiple users. This decreases the required connectivity bandwidth between players, and simplifies game creation by allowing standardized inter-game user context and action interfaces
- Touring: A user drives down the street, turns their head, and can see the URL or descriptive information associated with the real-world image.
- Military: By incorporating what is typically called "Situational Awareness" with first-person point of view displays, valuable information can be presented to the soldier (*e.g.*, indications of friends and foes, indications of hazardous and safe areas and situations, etc.).
- Surveyors: Longitude/latitude, plot boundaries, and other useful marks can be superimposed on the landscape.
- Emergency Crews: Police officers can be presented with task-related information without explicit commands on their part.
- Operating a Vehicle: The system can detect which car a user is driving. Perhaps the user has two cars; a sports car and a large van. They usually drive a sports car, but are currently driving the van. As they approach a curve in the road at a velocity that would be safe in the sports car, but is unsafe in the van, an indication of this situation can be displayed to the user in such a way so as not to excessively distract them and therefore cause a hazard. If they continue to approach the curve at an unsafe speed, the prominence and/or choice of output device can be modified to emphasize the hazard.
- Navigation: Integrated into the user's first person point of view is information useful in determining where they are, and how to follow their current route.

09636077-041304

As another example, an augmented reality tour of a physical attraction (such as a town hall or theme park) can be easily created with this system by creating a database of attraction records. Each record can contain real-world characterization information (*e.g.*, location, hours of operation, possible and optimal routes) and user-  
5 filtered interest characterizations (*e.g.*, type of attraction, amount of time required, recommendations of trusted reviewers, supplemental details on attraction or attraction elements, etc.). The database data can be made available to an interpretive engine, which by making use of the user's context model's location attributes can determine which attraction the user is closest to and can provide an indication of such without the user's  
10 explicit request. This represents a great simplification, since the creation of additional tour applications only requires the collection of new attraction information and the recording of it in a database. Not only can the interpretive engine be reused, but arbitrary changes in the user's context, including the availability of input/output devices, can be disregarded when creating additional tours.

Those skilled in the art will also appreciate that in some embodiments the functionality provided by the routines discussed above may be provided in alternate ways, such as being split among more routines or consolidated into less routines. Similarly, in some embodiments illustrated routines may provide more or less  
20 functionality than is described, such as when other illustrated routines instead lack or include such functionality respectively, or when the amount of functionality that is provided is altered. Those skilled in the art will also appreciate that the data structures discussed above may be structured in different manners, such as by having a single data structure split into multiple data structures or by having multiple data structures consolidated into a single data structure. Similarly, in some embodiments illustrated data  
25 structures may store more or less information than is described, such as when other illustrated data structures instead lack or include such information respectively, or when the amount or types of information that is stored is altered.

From the foregoing it will be appreciated that, although specific embodiments have been described herein for purposes of illustration, various  
30 modifications may be made without deviating from the spirit and scope of the invention.

